

Learning an Eddy Viscosity Model with Shrinkage– A Case Study with Jet-in-Crossflow Configuration

Jaideep Ray* and Sophia Lefantzi†
Sandia National Laboratories, Livermore, CA, 94550

Srinivasan Arunajatesan‡ and Lawrence Dechant§
Sandia National Laboratories, Albuquerque, NM, 87185

We demonstrate a statistical method for learning a high-order eddy viscosity model from experimental data and using it to improve the predictive skill of a Reynolds-averaged Navier-Stokes (RANS) simulator. The method is tested in a supersonic-jet-in-transonic-crossflow configuration. The process starts with a cubic eddy viscosity model (CEVM), calibrated for incompressible flows. It is fitted to measurements of turbulent stresses from a compressible flow experiment using shrinkage regression, specifically LASSO. LASSO retains the terms in the CEVM that are strongly supported by the data i.e., the most important terms, while removing the rest. For our particular case, LASSO removes all the terms except one that is quadratic in vorticity. The second step involves calibrating three parameters of the RANS model (one being the coefficient of the vorticity term in the truncated CEVM) using measurements of mean flow from a jet-in-crossflow experiment. The predictions of the calibrated RANS model with (truncated) CEVM is compared with experimental data as well as a calibrated RANS model using a linear eddy viscosity model (LEVM). Preliminary results show that high-order eddy viscosity model provides better predictions of turbulent stresses vis-à-vis RANS simulations with LEVM.

Nomenclature

c_3	=	Constant in the quadratic eddy viscosity model learnt via shrinkage
C_1, C_2	=	Constants in the equation for the evolution of ϵ
Ω^{obs}	=	Experimental observations of crossplane vorticity
$\Omega(C)$	=	RANS predictions of vorticity, for a specified (c_3, C_2, C_1)
η	=	Discrepancy between observed and modeled vorticity
$\mathcal{N}(\mu, \sigma^2)$	=	A normal distribution with mean μ and variance σ^2

I. Introduction

K- ϵ models, are routinely used in RANS simulations. However, in complex flow interactions, such as separated flows, flows with strong curvature etc. their predictive skills leave much to be desired. A jet-in-crossflow interaction is one such example, which is the subject of this study. The lack of predictive skill of RANS is caused by (1) approximations inherent in the k- ϵ model i.e., missing physics, also called structural or model-form error and (2) the use of suboptimal values for critical turbulence model parameter e.g., $C = (C_\mu, C_2, C_1)$. The nominal values of these parameters, $C_{\text{nominal}} = (0.09, 1.92, 1.44)$, are obtained by calibrating to canonical flows (e.g., shear and channel flows) that have little in common with jet-in-crossflow (JIC) interactions. In our previous paper¹ we addressed the problem of suboptimal values of C by performing a parameter calibration using experimental measurements of a supersonic-jet-in-transonic crossflow interaction^{2,3,4}. We posed a Bayesian inverse problem for C and developed a

* Technical Staff, Quant. Modeling & Analysis, MS 9159.

† Technical Staff, Quant. Modeling & Analysis, MS 9152, Senior Member AIAA.

‡ Technical Staff, Aerosciences Department, MS 0825, Senior Member AIAA.

§ Technical Staff, Aerosciences Department, MS 0825, Senior Member AIAA.

joint probability density function (PDF) using a Markov chain Monte Carlo (MCMC) sampler. In order to facilitate the $O(10^4)$ invocations of the forward model (the RANS) simulator that MCMC required, we replaced the RANS simulator with a surrogate model. Samples taken from the PDF were used to run an ensemble of RANS simulations to predict flow quantities which were not used in the inverse problem but for which we had experimental measurements. The Bayesian calibration improved the entire mean flowfield, but predictions of the turbulent stresses were not much better than those provided by C_{nominal} (note that it is the gradient of turbulent stresses that appear in the conservation laws, not the stresses themselves). The calibration effort showed that one of the primary sources of model-form error was the simplicity of the linear eddy viscosity model (LEVM) used in the RANS model. In this work, we address the remaining source of error in RANS models *viz.*, model-form error.

Enriching the LEVM with high-order terms can reduce the model-form error in RANS simulations; Craft⁵ describes a couple of quadratic and cubic extensions. While LEVM models the turbulent stress τ_{ij} using the turbulent kinetic energy k and the strain rate S_{ij} , the higher order terms contain the vorticity tensor Ω_{ij} as well. These terms are substantial in flows with strong curvature and/or vortical flows, such as JIC interactions. The high-order models have more model parameters - the quadratic models have 3 quadratic terms and the cubic model has 4 cubic terms, each with its scaling coefficient. These parameters are calibrated using simple incompressible flow interactions, such as curved channel flow, that may not be very relevant in transonic flows. This shortcoming can, in principle, be removed by calibration to experimental data. However, these models contain dissipation ϵ , which is not available outside a RANS k - ϵ simulator; thus the eddy viscosity model parameters can only be calibrated when integrated into a k - ϵ model.

Calibrating a high-order eddy viscosity model, in conjunction with a RANS simulator, is extremely challenging. The high-dimensional inverse problem (7 terms for a CEVM, plus any RANS-specific parameters) is computationally expensive for any except the simplest of flows. Constructing a surrogate model also does not offer a solution; properly sampling a 7-10 dimensional space, for the complex response surfaces that we might expect from a RANS JIC simulation would require a training set with $O(10^D)$ runs (D being the dimensionality of the parameter space). Further, given the limited experimental data, it may not be possible to estimate so many parameters with any degree of certainty. Thus dimensionality reduction is a necessity in this particular calibration problem.

Below we describe a two-step approach to performing the calibration. The first step is aimed at enriching the LEVM with as many high-order terms as can be supported by experimental data; each term (which contains S_{ij} and/or Ω_{ij}) has an unknown scaling coefficient. The second step involves implementing the enriched eddy viscosity model (EVM) into a k - ϵ RANS simulator and estimating the parameters of the enriched EVM as well as the k - ϵ model. Post-calibration, the predictions of the RANS simulator are compared with experimental data and predictions using a RANS simulator with a properly calibrated LEVM¹.

II. Technical Foundations

The primary challenge in this calibration exercise is the identification of a suitable model-form for the EVM, while balancing the confounding problems of limited experimental data, a large number of parameters and costly RANS simulations. In the first step we focus on estimating the model-form of the EVM from experimental data, in isolation from the RANS simulator *i.e.*, in the absence of ϵ that is required by the EVM. Experimental data consists of measurements of Reynolds stresses at a set of 315 locations (henceforth, probes) as well as PIV measurements of the mean flow (velocities) in the plane containing the probes *i.e.*, velocity gradients at the probes can be computed. The approach rests on two key hypotheses:

1. *Equilibrium in the production and dissipation of k* : A very crude, order-of-magnitude estimate of ϵ can be calculated by assuming that the rate of production of turbulent kinetic energy k is balanced by its dissipation; the flow quantities required to compute the production rate were measured in the experiment. Note that this assumption does not hold true everywhere and ϵ may assume non-physical (negative) values at certain probes.
2. *Discovery of a model-form via shrinkage regression*: In principle, a CEVM with 7 parameters could be calibrated using the crude estimate of ϵ at 315 probes. However, discarding probes with negative ϵ drastically reduces the available experimental data. The linearity of CEVM in the coefficients multiplying its terms allows us to pose a parameter estimation problem using LASSO⁶, a shrinkage method for linear

regression problems. LASSO only estimates coefficients in the CEVM that can be strongly inferred from the experimental data; it removes the other terms. We hypothesize that while the approximation errors in ϵ make the parameter estimates untrustworthy, the identification and retention of the most important terms in the CEVM are performed correctly.

II.A Learning a model-form for the high-order EVM

Let $\mathbf{V}^{(obs)}$ be observations of turbulent stresses at the set of probes where $\epsilon > 0$. Let $\mathbf{V}(\mathbf{c})$ be Craft's CEVM⁵, with $\mathbf{c} = \{c_i\}$, $i = 1 \dots 7$ being the vector of CEVM coefficients. We pose the estimation problem for \mathbf{c} as

$$(Eq. 1) \quad \min_{\mathbf{p}} \|\mathbf{V}^{(obs)} - \mathbf{V}(\mathbf{c})\|_2^2 + \lambda \|\mathbf{c}\|_1,$$

where $\|\cdot\|_2$ is the L2 norm (square root of the sum of squares), $\|\cdot\|_1$ is the L1 norm (sum of absolute values) and λ is a penalty. The minimization problem is solved using LASSO, which seeks to find parameters \mathbf{c} that reduce the data – model mismatch ($\mathbf{V}^{(obs)} - \mathbf{V}(\mathbf{c})$) as well as the penalty term $\|\mathbf{c}\|_1$. The simplest way to do so is to set as many c_i to 0 as possible, thus providing a *sparse* solution to the minimization problem. “Sparsifying” leads to the retention of the smallest set of high-order terms in the CEVM that can be justified by $\mathbf{V}^{(obs)}$.

The value of λ plays a crucial role in the robustness, simplicity and accuracy of the calibrated EVM. If the value of λ is too low, $\mathbf{V}(\mathbf{c})$ runs the danger of over-fitting $\mathbf{V}^{(obs)}$. This is quantified by cross-validation. $\mathbf{V}^{(obs)}$ is divided randomly into n equally sized bins (called “folds”). A value of λ is first set and we initially designate the first fold to be the “testing” fold. The EVM is learned (i.e., Eq. 1 is solved) on the remaining folds and its predictive accuracy is tested on the data in the “testing” fold, with the prediction error quantified as a root-mean-square-error (RMSE). This is repeated n times, by designating a new fold as the testing fold. The n RMSEs are summarized by their mean and standard deviations; they are characteristic of the chosen value of λ . The process is then repeated for different λ values. In case λ is too small, and most of the parameters (terms) in the CEVM are retained, the model may over-fit the limited training data i.e., the fit to $\mathbf{V}^{(obs)}$ appears to be very good, but the model's predictive skill in the “testing” fold is low. If the value of λ is too high, the CEVM is over-simplified (too many terms are removed), and the model has difficulty fitting the training data and predicting the testing fold, leading to high errors. For the optimal value of λ (λ_{min}) the mean and standard deviation of the RMSEs is minimized. There also exists λ_{1se} , where the mean RMSE is the same as the RMSE observed 1 standard deviation away from the mean RMSE for λ_{min} . The model corresponding to λ_{min} is the most accurate model, whereas one can use the model corresponding λ_{1se} if one desires a very sparse model with reasonable accuracy.

The process described above yields $\mathbf{c} = \{c_i\}$, $i = 1 \dots 7$ with c_i that could not be estimated confidently from $\mathbf{V}^{(obs)}$ set to zero. This provides us with a simplified/sparsified version of CEVM. In principle, we could use the values of \mathbf{c} so obtained as the calibrated values; however, the approximations used to compute ϵ (which is used in CEVM) raise questions about their accuracy. Therefore we simply use the non-zero values of c_i to select the CEVM terms to be retained in our final EVM. The data-driven approach to selecting the high-order terms also improves the chances of being able to infer their values from experimental data in a Bayesian calibration setting.

II.B Calibrating a RANS simulator with a high-order EVM

The model-form learnt in II.A is integrated into a compressible k- ϵ RANS simulator. The model equations being solved are described in Ref. 1. Let $\mathbf{c}_s = \{c_i\}$, $c_i \neq 0$, be the set of model coefficients in the sparsified form of the EVM. The (RANS) model parameters being estimated are $\mathbf{C} = \{\mathbf{c}_s, C_2, C_1\}$; C_2 and C_1 appear in the equation for the turbulent kinetic energy k . They are inferred from the experimental data in Ref. 4. The experimental setup is shown in Fig.1 (left). A Mach 3.73 jet is introduced from the bottom of the test section into a Mach 0.8 crossflow, flowing left to right. The jet curls over to the right and evolves into a counter-rotating vortex pair (CVP) as seen in Fig.1 (right). PIV measurements of mean flow velocities and turbulent stresses $\tau_{ij} = \langle u_i u_j \rangle$ are obtained on the midplane (plane of symmetry) and the crossplane (which slices through the CVP). The vorticity on the crossplane, inside a window \mathcal{W} that frames one of the vortices of the CVP, is chosen as the calibration variable. Let $\mathbf{\Omega}^{(obs)}$ be the streamwise vorticity observed at a set of locations (“probes”) in \mathcal{W} . Let the set of probes be \mathcal{P} . Let $\mathbf{\Omega}(\mathbf{C})$ be the

RANS prediction of the same corresponding to a parameter combination \mathbf{C} . We relate the two by $\boldsymbol{\Omega}^{(obs)} = \boldsymbol{\Omega}(\mathbf{C}) + \boldsymbol{\eta}$, $\boldsymbol{\eta} = \{\eta_j\}$, $\eta_j \sim \mathcal{N}(0, \sigma^2)$, $j \in \mathcal{P}$ i.e., the observed and predicted vorticity at a probe differ by η_j , a composite of measurement and model-form errors; η_j collected over \mathcal{P} resemble draws from a normal distribution with zero mean and an unknown standard deviation σ . A Bayesian inverse problem is formulated for the posterior distribution of \mathbf{C} , conditional on $\boldsymbol{\Omega}^{(obs)}$ $P(\mathbf{C}, \sigma | \boldsymbol{\Omega}^{(obs)})$

$$(Eq. 2) \quad P(\mathbf{C}, \sigma | \boldsymbol{\Omega}^{(obs)}) \propto \frac{1}{\sigma^{|\mathcal{P}|}} \exp\left(-\frac{\|\boldsymbol{\Omega}^{(obs)} - \boldsymbol{\Omega}(\mathbf{C})\|_2^2}{2\sigma^2}\right) \Pi(\mathbf{C}),$$

where $\Pi(\mathbf{C})$ is our prior belief regarding the where the appropriate values of \mathbf{C} might lie. The distribution $P(\mathbf{C}, \sigma | \boldsymbol{\Omega}^{(obs)})$ is arbitrary i.e., it does not belong to one of the analytical families such as Normal, Gamma etc. Consequently, we use a MCMC algorithm⁷, to draw (\mathbf{C}, σ) samples which are then histogrammed to visualize $P(\mathbf{C}, \sigma | \boldsymbol{\Omega}^{(obs)})$. The computational cost of generating $\boldsymbol{\Omega}(\mathbf{C})$ is addressed by developing polynomial surrogates of the RANS simulator. The process for constructing surrogates is described in Ref. 1.

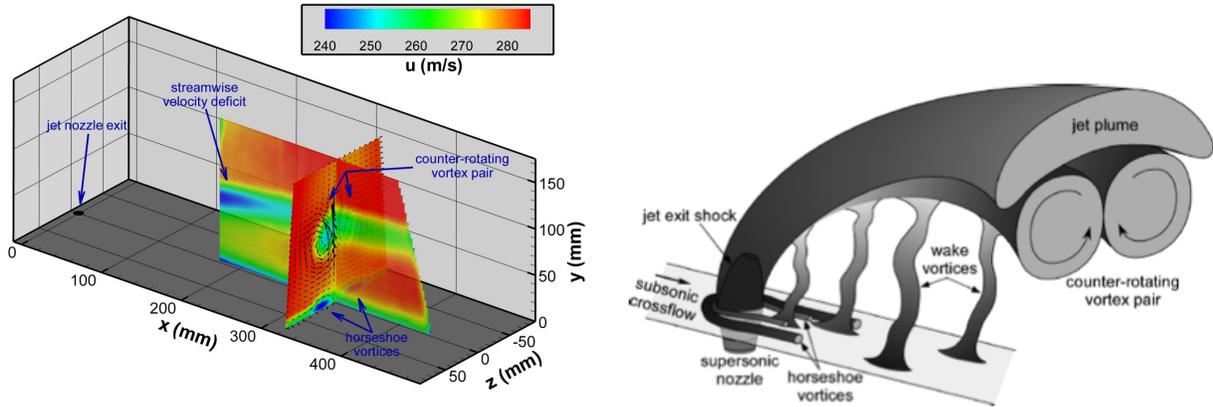


Fig. 1 Left: Schematic of the test section. The $M = 3.73$ jet is introduced from the bottom on the left. The crossflow, at $M = 0.8$, flows left to right. The midplane and crossplane where PIV measurements of mean flow velocities and turbulent stresses are available are show. The flow is statistically steady. Right: The jet, on interacting with the crossflow, curves and rolls into a counter-rotating vortex pair (CVP).

II.C Developing an informative prior distribution $\Pi(\mathbf{C})$

Upper and lower bounds on each of its constituents can delimit the parameter space for \mathbf{C} ; they form a hypercube \mathcal{H} . However \mathcal{H} is not very useful for calibration; many parameter combinations \mathbf{C} drawn from \mathcal{H} are unrealizable and either cause the RANS simulator to crash, not lead to a converged solution i.e., no steady-state flow solution exists, or provide flow solutions that do not resemble compressible flows. Consequently we construct a prior distribution $\Pi(\mathbf{C})$ in the manner described in Ref. 1. The parameter space is sampled (we refer to this as the “training set”) and RANS simulations are initiated. Of the runs that were realized, we compute the mean square error (MSE) $\|\boldsymbol{\Omega}^{(obs)} - \boldsymbol{\Omega}(\mathbf{C})\|_2^2$. We choose the runs whose MSE lie below the first quartile. The parameters \mathbf{C}_i that correspond to these “good runs” describe a coherent region \mathcal{R} in \mathcal{H} where we believe that the optimal parameters may be found. We define our prior belief $\Pi(\mathbf{C})$ as being 1 for $\mathbf{C} \in \mathcal{R}$ and 0 otherwise.

In order to formalize the definition of \mathcal{R} so that we can include it in the MCMC calibration, we construct a classifier. The process is described in Ref. 1. We define a binary variable $\zeta(\mathbf{C})$, such that $\zeta(\mathbf{C}) = 1$, for all \mathbf{C} corresponding to the “chosen runs”; else $\zeta(\mathbf{C}) = 0$. We train a support vector machine classifier (SVMC) on this dataset, and incorporate it into the right hand side of Eq. 2 to serve as $\Pi(\mathbf{C})$ i.e., given an arbitrary \mathbf{C} , the classifier identifies whether it lies in \mathcal{R} ($\Pi(\mathbf{C}) = 1$) or not ($\Pi(\mathbf{C}) = 0$).

III. Calibration Results

A model-form for the EVM: Turbulent stress measurements on the midplane are used to learn the model form for the EVM. As described in Sec.II.A, we compute an approximate ε and discard probes where $\varepsilon < 0$. Only 103/315 probes survive the cut. At each probe, we employ three observables τ_{11} , τ_{22} and τ_{12} i.e., we have 306 observations. The shrinkage is initiated with Craft's CEVM (which has 7 parameters). Note that while 306 data points may seem sufficient to infer 7 parameters, there is no guarantee that the midplane data are actually informative about all the CEVM parameters. In Fig 2 (left) we plot the performance of the shrinkage for different values of λ . Deviance is the difference between the $\mathbf{V}^{(\text{obs})}$ and $\mathbf{V}(\mathbf{c})$ when all high-order terms are removed. We see that as λ increases, the ability of the model to fit the training data and explain the deviance decreases. Note that no cross-validation was performed. In Fig. 2 (mid) we see how increasing λ reduces the number of non-zero parameters in the CEVM. In Fig. 2 (right), we see the results from the 15-fold cross-validation performed for various values of $\ln(\lambda)$. The values of λ_{\min} and $\lambda_{1\text{se}}$ are also plotted, along with the complexity (the number of non-zero c_i) of the CEVM. Note that for $\lambda_{1\text{se}}$ only one term is retained in the EVM; the term is quadratic in vorticity. Given the strongly vortical nature of the flow, this is not entirely surprising. In Table 1 we tabulate c_i obtained using λ_{\min} and $\lambda_{1\text{se}}$; they are quite different from c_i obtained by Craft from incompressible curved channel flows. Given the simple enhancement of the LEVM that $\lambda_{1\text{se}}$ provides, we will proceed with a LEVM enriched with a quadratic vortical term and $\mathbf{C} = \{c_3, C_2, C_1\}$. We will refer to it as the quadratic eddy viscosity model, QEVM.

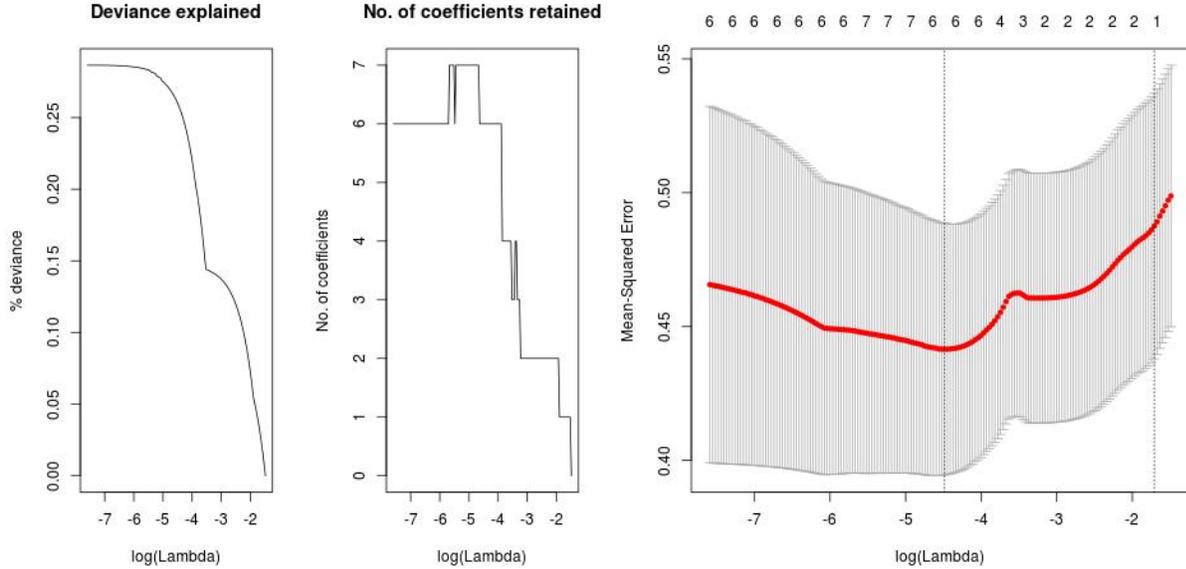


Fig 2. Left: The reduction in fraction of the deviance explained as the EVM model is simplified by penalizing the number of terms it contains i.e., by increasing λ . Middle: The reduction in the number of model parameters in the EVM model as λ is increased. Right: Results from the 15-fold cross-validation of the CEVM model. The red line is the mean RMSE whereas the error bars are the standard deviations. The two vertical lines are λ_{\min} and $\lambda_{1\text{se}}$.

Constructing the prior $\Pi(\mathbf{C})$: The first step in the construction of the informative prior is to explore the 3-dimensional parameter space \mathcal{H} for $\mathbf{C} = \{c_3, C_2, C_1\}$. The bounds were: $0 \leq c_3 \leq 3.5$, $1.7 \leq C_2 \leq 2.5$, $1.2 \leq C_1 \leq 1.7$. The bounds for C_2 and C_1 were obtained from Ref. 8. 2744 $\{c_3, C_2, C_1\}$ samples were selected from \mathcal{H} using a quasi-Monte Carlo, space-filling Halton sequence. Of these, 222 runs completed, which are plotted in Fig. 3 (left). We clearly see that lower values of c_3 are favored, though smaller values of C_2 and C_1 may allow c_3 to assume values around 0.6. No run with a value of $c_3 > 1$ was realized, showing the degree of error introduced into Table 1 due to the use of approximate values of ε in the CEVM. The selection of the top quartile of runs did not allow the training of an accurate SVMC (misclassification rate $< 10\%$) due to the paucity of training data, though it did allow us to select a more relevant training set of 1500 \mathbf{C} samples from the populated portion of \mathcal{H} , as seen in Fig. 3 (left). The training data is still being generated.

Method	c_1	c_2	c_3	c_4	c_5	c_6	c_7	MSE
Craft	-0.1	0.1	0.26	-10	0	-5	5	0.662
λ_{\min}	-0.052	-0.061	1.56	-2.36	7.1	4.27	0	0.386
λ_{1se}	0.0	0.0	0.397	0.0	0	0	0	0.483

Table 1: Parameters of the CEVM estimated using λ_{\min} and λ_{1se} , compared with the classical values obtained from incompressible curved channel flows by Craft.

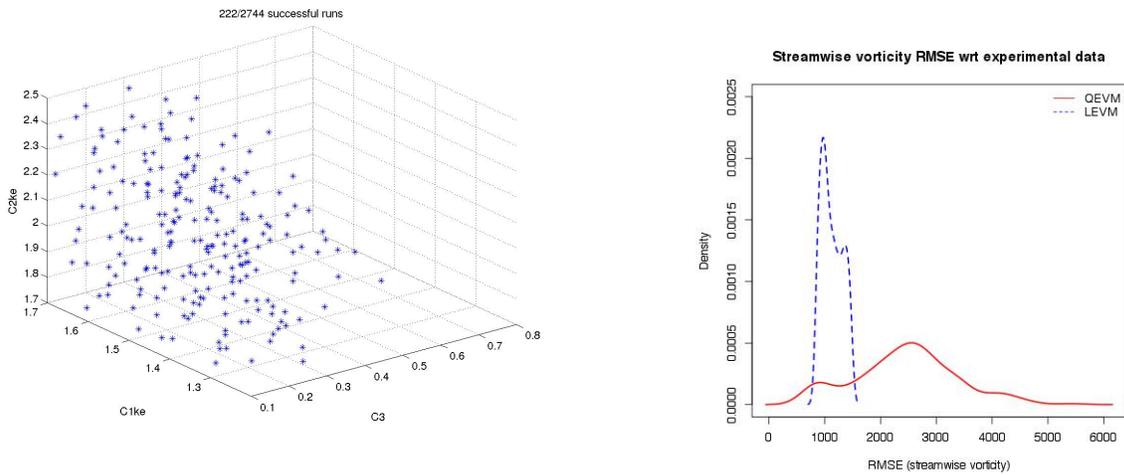


Fig. 3: Left: Distribution of \mathbf{C} in \mathcal{H} corresponding to the 222 QEVM runs that successfully completed. It shows the realizable part of \mathcal{H} . Right: Distributions of RMSEs of streamwise vorticity, vis-à-vis experimentally observed values. The two distributions correspond to simulations with LEVM and QEVM.

Calibrating a RANS simulator with the QEVM: In the absence of a proper training set, we have not yet been able to perform a Bayesian calibration for $\mathbf{C} = \{c_3, c_2, c_1\}$. However, the 222 successful simulations do allow us to perform a preliminary test to check if the QEVM has the potential to provide better predictions of vorticity on the crossplane compared to a calibrated LEVM and determine the nature of QEVM predictions using best QEVM run. The calibration of the LEVM is described in Ref. 1. In Fig. 3 (right) we plot the distribution of RMSE in the crossplane vorticity (vis-à-vis experimental values) from the top quartile of simulations performed using LEVM and the 222 QEVM runs that were successful. We see from the PDF that QEVM RMSEs can assume smaller values than LEVM RMSEs. The distribution of errors allows us to select the best performing QEVM and compare it's predictions with those obtained from a calibrated LEVM.

In Fig. 4, we plot the predictions of τ_{11} (first column), τ_{22} (second column) and τ_{12} (third column) on the midplane computed using the calibrated LEVM, the best-case QEVM and the experimental measurements. These are plotted at 3 distances downstream of the jet – 200 mm (first row), 250 mm (second row) and 300 mm (bottom row). It is quite clear that the QEVM performs far better than LEVM. QEVM (note: the parameter set is *not* calibrated to the data; we simply chose the best run of the 222 that were realized) tends to over-predict τ_{22} whereas LEVM underpredicts it; as one proceeds downstream, QEVM provides better predictions of τ_{22} than LEVM. In contrast LEVM predicts τ_{11} better, with QEVM overpredicting it consistently.

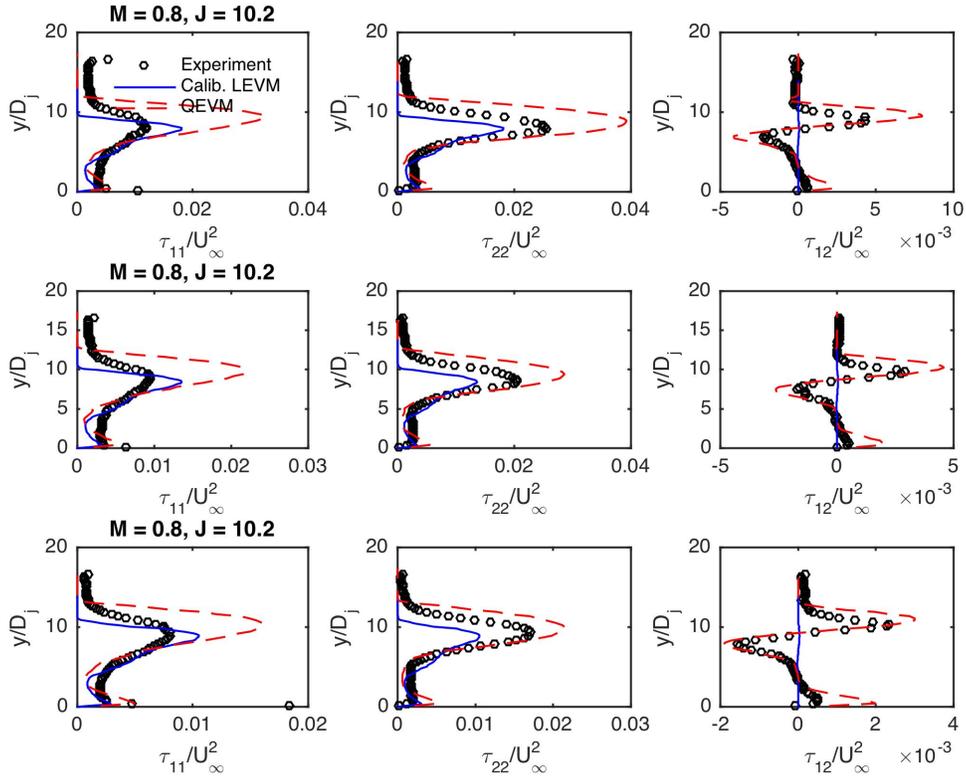


Fig. 4: Comparison of Reynolds stresses as predicted by a calibrated LEVM (blue line) and our best QEVM run (red line). The three rows of figures correspond to three station 200mm, 250mm and 300mm downstream of the jet. The three columns are plots of τ_{11} , τ_{22} and τ_{12} respectively. QEVM's predictions of τ_{12} are quite impressive.

IV. Conclusion, Outline of the Final Paper and Future Work

We have developed a method for addressing the model-form errors in RANS simulations. The method seeks to learn a better form for the eddy viscosity model from experimental data using LASSO. Due to the approximations inherent in how the model discovery problem is formulated, we only accept the model form that is discovered and ignore the model parameters that are also estimated by LASSO. Our method revealed that given the data at hand, the linear eddy viscosity model could be enriched with a quadratic term in vorticity (for our strongly vortical flow); further, the corresponding model parameter could be estimated from data. We refer to the enriched EVM as the QEVM. A RANS simulator with the QEVM has been constructed and is being used to generate a training dataset. Preliminary runs show that the QEVM has the potential to be more accurate and is being used to generate a training dataset. Preliminary runs show that the QEVM has the potential to be more accurate and is being used to generate a training dataset. Preliminary runs show that the QEVM has the potential to be more accurate and is being used to generate a training dataset.

In the final paper, we will construct surrogate models of vorticity predictions on the crossplane by RANS-QEVM. These surrogates will be used in a Bayesian calibration of \mathbf{C} to construct the posterior distribution $P(\mathbf{C}, \sigma \mid \mathbf{\Omega}^{(obs)})$. The predictive skill of $P(\mathbf{C}, \sigma \mid \mathbf{\Omega}^{(obs)})$ will be checked via posterior predictive tests for the mean flow on the midplane. The predictions will be compared with LEVM results from our previous paper¹.

Acknowledgments

This work was supported by Sandia National Laboratories' Advanced Scientific Computing (ASC) Verification

and Validation program. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U. S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

¹J. Ray, S. Lefantzi, S. Arunajatesan and L. Dechant, "Bayesian calibration of a RANS model with a complex response surface - A case study with jet-in-crossflow configuration", AIAA-2015-2784, *45th AIAA Fluid Dynamics Conference, Dallas, TX*, June 22-26, 2015.

²S. J. Beresh, J. F. Henfling, R. J. Erven and R. W. Spillers, "Penetration of a transverse supersonic jet into a subsonic compressible crossflow", *AIAA Journal*, vol. 43, no. 2, pp. 379-389, 2005.

S. J. Beresh, J. F. Henfling, R. J. Erven and R. W. Spillers, "Turbulent characteristics of a transverse supersonic jet in a subsonic compressible crossflow", *AIAA Journal*, vol. 43, no. 11, pp. 2385-2389, 2005

⁴S. J. Beresh, J. F. Henfling, R. J. Erven, and R. W. Spillers, "Crossplane velocimetry of a transverse supersonic jet in a transonic crossflow," *AIAA Journal*, vol. 44, no. 12, pp. 3051–3061, 2006.

⁵T. J. Craft, B. E. Launder and K. Suga, "Development and application of a cubic eddy-viscosity model of turbulence," *International Journal of Heat and Fluid Flow*, vol. 17, pp. 108-115, 1995.

⁶T. Hastie, R. Tibshirani and J. Friedman, "The elements of statistical learning", Springer, 2nd Edition, 2008.

⁷H. Haario, M. Laine, A. Mira and E. Saksman, "DRAM: Efficient adaptive MCMC," *Statistical Computing*, vol. 16, pp. 339-354, 2006.

⁸S. Guillas, N. Glover and L. Malki-Epshtein, "Bayesian calibration of the constants of a k- ϵ turbulence model for a CFD model of street canyon flow", *Computational Methods in Applied Mechanical Engineering*, vol. 279, pp. 536-553, 2014.